

# Common Reasons Why Jobs Won't Start

If your job does not run after it has been successfully submitted, it might be due to one of the following reasons:

- The queue has reached its maximum run limit
- Your job is waiting for resources
- Your mission share has run out
- The system is going into dedicated time
- Scheduling is turned off
- Your job has been placed on hold
- Your home filesystem or default /nobackup filesystem is down
- Your hard quota limit for disk space has been exceeded

Each scenario is described in the sections below, along with troubleshooting steps. The following commands provide useful information about the job that can help you resolve the issue:

`tracejob job_id`

Displays log messages for a PBS job on the PBS server. For Pleiades, Aitken, and Electra jobs, log into pbspl1 before running the command. For Endeavour jobs, log into pbspl4 before running the command.

`qstat -as job_id`

Displays all information about a job on any Pleiades front end (PFE) node. For Endeavour jobs, the *job\_id* must include both the job sequence and the server name, pbspl4. For example:

```
qstat -as 2468.pbspl4
```

## The Queue Has Reached Its Maximum Run Limit

Some queues have a maximum run (`max_run`) limit—a specified maximum number of jobs that each user can run. If the number of jobs you are running has reached that limit, any jobs waiting in the queue will not begin until one of the running jobs is completed or terminated.

Currently, the `debug` and `ldan` queues each have a `max_run` limit of 2 per user.

## Your Job is Waiting for Resources

Your job might be waiting for resources for one of the following reasons:

- All resources are busy running jobs, and no other jobs can be started until resources become available again
- PBS is draining the system and not running any new jobs in order to accommodate a high-priority job
- Users have submitted too many jobs at once (for example, more than 100), so the PBS scheduler is busy sorting jobs and cannot start new jobs effectively
- You requested a specific node or group of nodes to run your job, which might cause the job to wait in the queue longer than if nodes were not specified

To view job status and events, run the `tracejob` utility on the appropriate PBS server (pbspl1 for Pleiades, Aitken, and Electra jobs, or pbspl4 for Endeavour jobs). For example:

```
pfe26% ssh pbspl1
pbspl1% tracejob 234567
```

Job: 234567.pbspl1.nas.nasa.gov

```
02/15/2019 00:23:55 S Job Modified at request of Scheduler@pbspl1.nas.nasa.gov
02/15/2019 00:23:55 L No available resources on nodes
02/15/2019 00:38:47 S Job Modified at request of Scheduler@pbspl1.nas.nasa.gov
02/15/2019 00:50:30 S Job Modified at request of Scheduler@pbspl1.nas.nasa.gov
02/15/2019 01:51:21 S Job Modified at request of Scheduler@pbspl1.nas.nasa.gov
02/15/2019 01:55:38 S Job Modified at request of Scheduler@pbspl1.nas.nasa.gov
02/15/2019 02:16:44 S Job Modified at request of Scheduler@pbspl1.nas.nasa.gov
02/15/2019 07:39:48 L Considering job to run
02/15/2019 07:39:48 L Job is requesting an exclusive node and node is in use
```

TIP: If the scheduler has not yet reviewed the job, no information will be available and the **tracejob** utility will not provide any output.

If your job requests an exclusive node and that node is in use, you can wait for the requested node, request a different node, or submit the job again without requesting a specific node.

To view the current node usage for each processor type, run the **node\_stats.sh** script. For example:

```
pfe26% node_stats.sh
```

Node summary according to PBS:

Nodes Available on Pleiades	: 10444	cores: 225180
Nodes Available on Aitken	: 2146	cores: 170296
Nodes Available on Electra	: 3338	cores: 120344
Nodes Down Across NAS	: 1311	

Nodes used/free by hardware type:

Intel SandyBridge	cores/node:(16)	Total: 1683, Used: 1527, Free: 156
Intel IvyBridge	(20)	Total: 4809, Used: 4578, Free: 231
Intel Haswell	(24)	Total: 1969, Used: 1891, Free: 78
Intel Broadwell	(28)	Total: 1924, Used: 1820, Free: 104
Intel Broadwell (Electra)	(28)	Total: 1098, Used: 1037, Free: 61
Intel Skylake (Electra)	(40)	Total: 2240, Used: 2200, Free: 40
Intel Cascadelake (Aitken)	(40)	Total: 1099, Used: 1022, Free: 77
AMD ROME EPYC 7742 (Aitken)	(128)	Total: 987, Used: 923, Free: 64

Nodes currently allocated to the gpu queue:

Sandybridge (Nvidia K80)	Total: 59, Used: 1, Free: 58
Skylake (Nvidia V100)	Total: 19, Used: 5, Free: 14
Cascadelake (Nvidia V100)	Total: 36, Used: 30, Free: 6

Nodes currently allocated to the devel queue:

SandyBridge	Total: 381, Used: 301, Free: 80
IvyBridge	Total: 318, Used: 201, Free: 117
Haswell	Total: 142, Used: 66, Free: 76
Broadwell	Total: 146, Used: 79, Free: 67
Electra (Broadwell)	Total: 0, Used: 0, Free: 0
Electra (Skylake)	Total: 0, Used: 0, Free: 0
Aitken (Cascadelake)	Total: 0, Used: 0, Free: 0
Aitken (Rome)	Total: 0, Used: 0, Free: 0
Skylake gpus	Total: 1, Used: 0, Free: 1
Cascadelake gpus	Total: 0, Used: 0, Free: 0

Jobs on Pleiades are:

```
requesting: 42 SandyBridge, 5799 IvyBridge, 4938 Haswell, 1521 Broadwell,
158 Electra (B), 1600 Electra (S), 782 Aitken (C), 460 Aitken (R) nodes
using: 1527 SandyBridge, 4578 IvyBridge, 1891 Haswell, 1820 Broadwell,
1037 Electra (B), 2200 Electra (S), 1022 Aitken (C), 923 Aitken (R) nodes
```

TIP: Add **/u/scicon/tools/bin** to your path in **.cshrc** or other shell start-up scripts to avoid having to enter the complete path for this tool on the command line.

## Your Mission Share Has Run Out

If all of the cores within your mission directorate's share have been assigned or if running the new job would exceed your mission share, your job will not run. If resources appear to be available, they belong to other missions.

To view all information about your job, run `qstat -as job_id`. In the following sample output, a comment (line 5) indicates that the job would exceed the mission share:

```
pfe21% qstat -as 778574
JobID      User      Queue    Jobname   CPUs   wallt   Ss   wallt   Eff   wallt
-----
778574.pbspl1  zsmith   normal   my_GC     12    04:00   Q    07:06   --    04:00
Job would exceed mission CPU share
```

To view the distribution of shares among all mission directorates, run `qstat -W shares=-`. For example:

```
pfe21% qstat -W shares=-
Group  Share% Use%  Share Exempt    Use  Avail Borrowed Ratio Waiting
-----
Overall    100    1 502872      0  9620 493252      0  0.02    400
  ARMD     26    27 128245      0 135876      0  7631  1.06 579173
  HEOMD     22    22 108514    1832 111840      0  3326  1.03 1783072
  SMD       50    40 246626      32 201220  45406      0  0.82  94316
  NAS        2     0   9864     180  2416   7448      0  0.24    128
```

If running your job would exceed your mission share, you might be able to borrow nodes from other mission directorates. To borrow nodes, your job must not request a wall-clock time that is too long (more than 4 hours). See [Mission Shares Policy on Pleiades](#) for more information.

## The System is Going into Dedicated Time

When dedicated time is scheduled for hardware and/or software work, the PBS scheduler will not start a job if its projected end-time runs past the beginning of the dedicated time.

If you can reduce the requested wall-clock time so that your job will finish running prior to dedicated time, then your job can then be considered for running. To change the wall-clock time request for your job, follow the example below:

```
pfe21% qalter -l walltime=hh:mm:ss job_id
```

TIP: For Endeavour jobs, the *job\_id* must include both the job sequence number and the Endeavour server name (for example, 2468.pbspl4).

To find out whether the system is in dedicated time, run the `schedule all` command. For example:

```
pfe21% schedule all
No scheduled downtime for the specified period.
```

## Scheduling is Turned Off

Sometimes job scheduling is turned off by a system administrator. This is usually done when system or PBS issues need to be resolved before jobs can be scheduled to run. When this happens, you should see the following message near the beginning of the `qstat -au your_userid` output.

```
+++Scheduling turned off.
```

## Your Job Has Been Placed On Hold ("H" Mode)

A job can be placed on hold either by the job owner or by someone who has root privilege, such as a system administrator. If your job has been placed on hold by a system administrator, you

should get an email explaining the reason for the hold.

## **Your Home Filesystem or Default /nobackup Filesystem is Down**

When a PBS job starts, the PBS prologue checks to determine whether your home filesystem and default /nobackup filesystem are available before executing the commands in your script. If your default /nobackup filesystem is down, PBS cannot run your job and will put the job back in the queue. If your PBS job does not need any file in that filesystem, you can tell PBS that your job will not use the default /nobackup to allow your job to run.

For example, if your default is /nobackupp1, you can add the following in your PBS script:

```
#PBS -l nobackupp1=0
```

## **Your Hard Quota Limit for Disk Space Has Been Exceeded**

NAS quotas have hard limits and soft limits. Hard limits should never be exceeded. Soft limits can be exceeded temporarily, for a grace period of 14 days. If your data remains over the soft limit for more than 14 days, the soft limit is enforced as a hard limit.

See [Quota Policy on Disk Space and Files](#) for more information.

---

Article ID: 198

Last updated: 14 Oct, 2021

Revision: 47

Running Jobs with PBS -> Optimizing/Troubleshooting -> Common Reasons Why Jobs Won't Start

<https://www.nas.nasa.gov/hecc/support/kb/entry/198/>